RSRE MEMORANDUM No. 4277

ROYAL SIGNALS & RADAR ESTABLISHMENT

ALVEY MMI007: OBJECT IDENTIFICATION FROM 2D IMAGES: THE RSRE CONTRIBUTION

Authors: P Fretwell, R Series

PROCUREMENT EXECUTIVE,
MINISTRY OF DEFENCE,
R S R E MALVERN,
WORCS.

0046103

CONDITIONS OF RELEASE

BR-110653

****

U

****

Y

Reports quoted are not necessarily available to members of the public or to commercial organisations.

ROYAL SIGNALS AND RADAR ESTABLISHMENT

Memorandum 4277

TITLE: ALVEY MMI007: OBJECT IDENTIFICATION FROM 2D IMAGES: THE RSRE CONTRIBUTION

AUTHORS: P. Fretwell and R. Series

DATE: MAY, 1989

SUMMARY

The Alvey program was a government scheme to improve the United Kingdoms competitiveness in the field of information technology. It was coordinated by the Alvey directorate. The directorate was manned from the SERC, MoD, the academic world and people from industry. Alvey encouraged the universities and industry to work together by helping to fund proposals for pre-competitive research on a series of strategic areas in information technology.

A consortia was formed in 1984 to tackle the subject of automatically finding and identifying specified objects in a complicated environment from two dimensional images. The consortia members consist of four industrial partners - British Aerospace, Marconi Command and Control Systems, Standard Telecommunications Laboratories and The Royal Signals and Radar Establishment - and three universities - Bristol, Reading and Surrey. The project began in 1985 and now in 1989 is almost completed. This report describes the RSRE contribution to the MMI007 project.

ALVEY

ALVEY PROJECT No. MMI C07

OBJECT IDENTIFICATION FROM 2D IMAGES

○ British Aerospace ○ Marconi

○ STC ○ RSRE ○ Bristol University

○ Reading University ○ Surrey University

○ Rutherford Appleton Laboratory

A consortium was formed in 1984 to tackle the subject of 'Object Identification in 2D images'. The consortium consists of 4 industrial partners — British Aerospace, MCCS, STC and RSRE — and 3 universities — Bristol, Reading and Surrey. The project (MMI 007) began in January 1985 and is scheduled to run for three years; the final product being a software demonstrator of a working vision system.

Aims

The work is aimed at the problem of automatically finding and identifying specified objects in a complicated environment from a two-dimensional image. Complications may arise from several sources; such as the nature of the background, environmental variability, object/sensor geometry, object occlusion, distortion and noise. The research puts emphasis on establishing methods of general application rather than specific short-term solutions. The overall aims of the project are:-

- To discover and develop useful techniques for image analysis and

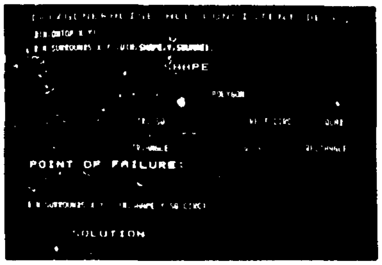

interpretation, including both low-level processing (segmentation and feature analysis) and high-level processing (involving knowledge, inference, learning and decision making).
- To show how successful these techniques are, either singly or combined by utilising application-oriented software.

Achievements

Two exemplars have been defined. The first is a small demonstration of resistor classification by the automatic identification of relevant colour codes. This was undertaken by STC. The second and main exemplar has been defined as the task of locating and identifying cars in outdoor environments, and is currently being tackled by all the consortium members. Many images have been collected and digitised. The data has then been processed using a number of algorithms developed during the project. These include:-

- *Edge finding and linking to form boundary descriptions*
- *Rule-based feature/primitive extraction*
- *Contextual classification/relaxation*
- *Colour segmentation*
- *Region labelling*
- *Rule-based interpretation*
- *Learning by rule induction*
- *Model matching*

Linking of these modules is now under way. A computer framework based on a blackboard is currently being developed so that the algorithms can be run under 'intelligent' control.

Future

Although there is further work to be done on some techniques, notably model matching, rule-based interpretation and learning, the emphasis is now on assessing the algorithms developed and improving the vision framework. By January 1988 the algorithms should be working together in a single vision system. Some work will still continue after this date on refinement of some of the high-level techniques. An image capture facility will also be set up at Reading University to provide controlled material for other application domains, such as part inspection and assembly. This facility will be available (within constraints) to other consortia.

for information contact:-

M.B. Brown

British Aerospace plc, Sowerby Research Centre (FPC 267)

PO Box 5, Filton, Bristol BS12 7QW, England

Telephone (0272) 693831, Ext. 6216

Contents

1 Introduction

At RSRE, participation in Alvey MMI007 has been through the Pattern Matching and Machine Intelligence division. The division, which receives support from DTI, MOD, Esprit and private industry, is devoted to developing techniques for pattern recognition and machine intelligence. RSRE has contributed to the collaboration through the development of advanced edge detection algorithms, extraction of object cues, and investigation of object representations.

Extraction of reliable edge information from an image is of fundamental importance in many image processing applications. Despite the apparent simplicity, successful edge delineation is a strong function of image content. In areas of high detail, retention of edge information requires the use of a small window, while for coarser features a large window is of value so as to minimise the effects of noise in the image. To address these and other issues RSRE has developed an rule based approach to edge detection. The algorithm continuously monitors the nature of the edge and optimises it's performance according to image content.

For many applications speed of edge detection is an important issue and the overheads of the rule based approach to edge detection cannot be afforded. RSRE has developed a second edge finder which is based on the Sobel operator. In tests the algorithm performed significantly faster than the Alvey Canney routines with no noticeable degradation in performance.

Detection of edges provides only the first step in object recognition. Many of the later algorithms have considerable computational cost and development of interest operators and filtering schemes are of importance. To this end RSRE has developed two routines. The first provides a classification of lines according to their fractal dimension. The routine proves useful for discrimination between edges due to foliage or vegetation and many man made features. The algorithm is fast and is also well suited to parallel processor arrays.

The second routine breaks the image into regions and analyses the statistical distribution of edge directions to find the dominant direction. This provides a useful cue to focus the search by more expensive operators.

Object representation is a major issue in image interpretation. While spatial information provides a valid means of describing specific instances of objects, there are many situations in which it would be desirable to be able to specify a class of objects by functional class as distinct to spatial appearance. A part of the contribution from RSRE has been to investigate methods for encapsulation of functional as well as spatial models. Work has been carried out to enhance the inference mechanism of PROLOG and to utilise heuristic knowledge derived by low level operators during the model matching process. Finally, in any reasoning system, uncertainty plays a major role. The potential use of 'FRIL', a fuzzy logic language, closely similar to 'PROLOG' has been evaluated.

2 RSRE's role in MMI007

Within the framework of the MMI007 Consortium, RSRE pursued research into techniques for extracting edge features from images, and using these to infer object cues which could be used to control the search space of top-down model driven methods developed by other partners.

The work had three principal aims:

i) Improving techniques for edge extraction, especially aimed at avoiding distortions and fragmentation which occurs with traditional edge finding methods.

ii) Developing intermediate level techniques able to take edge features and produce both generic and object specific cues.

iii) Review methods to accumulate evidence from collections of extracted cues and reason about object hypothesis. This had two elements, one based occurrence and simple relations between cues, and a more speculative activity aimed at using verbalised descriptions of object to generate structural models in terms of a hierarchy of relations with cues as terminal nodes.

Aims i) and ii) has led to techniques which have been used within the consortium as part of the integrated project, whilst iii), which has drawn upon expertise within Bristol University, has produced a number of individual techniques which require further work before demonstration is possible.

In addition to the Consortium activity, RSRE supported a related contract with Bristol University on the use of Conceptual Graphs in object inference, which was formulated as a direct extension to Aim (iii).

3 Technical Description of Work Carried out at RSRE

3.1 Edge Detection

3.1.1 Extraction of Boundaries using Local Measures Driven by Rules

This section outlines the rule based boundary extraction algorithm which is described in more detail in Sleigh (1986).

Boundary extraction is often performed by local edge finding algorithms such as Sobel (or Marr-Hildreth, (Marr, Hildreth, 1980) which give an analytical measure of the "edginess" of a particular region. These methods suffer from drawbacks associated with questions of scale, interaction between nearby image features and restriction to step edges which limits their use to images of good quality and sparse feature density. Various attempts have been made to produce more robust boundary extractors, including the scale-space methods of Witkin, (Witkin, 1983), the optimisation approach of Canney (Canney, 1983), and the application of weak-continuity of Blake (Blake 1983). All of these attempt to locally adjust the boundary finding process in response to the image data present, using weak heuristic knowledge which express some constraint imposed by the nature of the world.

In an attempt to rectify some of the deficiencies of existing edge detection algorithms, work at RSRE was initiated to develop a rule based approach. The approach identifies an edge as the boundary between two domains and is concerned with the application of strong boundary related constraints, but applied without recourse to object specific information. This contrasts with other methods for object detection such as are used by conventional edge extraction routines, or shape specific operators such as the generalised Hough transform. The method has been developed with the aim of robustly finding and discriminating boundaries of different types in close proximity, including boundaries which differ only in texture. The method consists of a set of "measurement probes" which evaluate local image properties such as gradient, roughness and autocorrelation, which are called and controlled by a system of rules which express our knowledge of the makeup of different types of edge. This knowledge is partly derived by common sense and partly from interactive training using example image sets.

The method recognises the main aim of boundary tracing: to accurately follow structural features in images. It succeeds in classifying and tracing difficult boundaries, and can cope with finely spaced features as well as simple texture boundaries, even when the feature strength is comparable with noise. This is achieved without object or shape specific prior knowledge. This involves general heuristic which cannot easily be expressed, and requires a richer classification of boundary type than has previously achieved. In addition, explicit account has been taken of complex edges which behave differently from the step edges assumed by many boundary extraction methods. The rules described are used to form plausible boundaries taking account of the properties adjacent to a boundary, and offer a general framework into which additional knowledge, such as texture description, can be added.

3.1.2 Advanced Sobel Edge Operator

This section describes briefly the edge operator described more fully in (Radford, 1989).

Although the rule based approach to edge detection potentially offers many advantages, these advantages are gained at the expense of computational speed. To provide a faster algorithm, RSRE has developed a fast edge-line extraction operator and fractal line discriminant operator. Although developed as separate routines, the two operators are intended to be used together.

The edge operator has been developed to run quickly on a serial machine although, being principally based upon near neighbour operations will also port efficiently to a SIMD parallel machine

The principal components of the edge detection algorithm are:

1. edge convolution,

2. resolution of edgel direction from convolution components,

3. thinning of edges to single pixel lines using non maximal suppression,

4. restriction of maximum angular deviation allowed along edge line,

5. application of hysteresis thresholding operation.

The operator is arranged to output three images: edge magnitude, direction, and a logical mask representation of the hysteresis thresholded edge map. To conserve space the magnitude and directional information are generally packed into a single 8 bit representation.

The operator also performs better than the standard Sobel operator on the extraction of fine detail especially on the extraction of corners (Booth, 1988). Unlike the Canny operator, smoothing applied by the Sobel is fixed and small, so the operator performs less well than the Canny in the presence of increasing noise. However, in practice this is rarely a problem as, in common with the majority of visible imagery, noise in the Alvey test images is generally originates from structure in the scene rather than camera noise.

The operator executes relatively fast. A typical 256x256 Alvey car image is processed in about 25 seconds on a VAX 11/780, some three times faster than for the Alvey Canny operator.

3.1.3 Fractal Line Discriminant Operator

This section describes briefly the fractal discriminant operator. The operator is described more fully in (Radford, 1989).

Inspection of edge extracted images from the Alvey test set shows a large percentage of the edge map originates from detail in foliage and vegetation in the image. In contrast, edge maps of many man made objects are generally straighter and smoother then those observed in foliage. When attempting to analyse such output by either symbolic or other methods (such as a Hough transform), the edge detail present in the foliage greatly increases the computational burden, and often obscures the features which are sought. The aim of this operator is to provide a discriminant against lines containing a lot of detail to ease the burden in the subsequent processing stages.

The filter may be thought of as approximating to a fractal analysis of the lines in the image. A straight line has a fractal dimension of unity which increases as the line becomes increasingly contorted. Although the simple discriminating measure used here is not rigorously related to the fractal dimension, in common with the fractal dimension of a true fractal curve, it is found to be independent of scale (magnification), over a reasonable range. This makes the operator useful even when the approximate image content is unknown.

The measure used is derived by summing the amount of turn between edgel elements along the lines in the image. If the total normalised angular turn along a line exceeds some threshold the line is rejected. The implementation of the operator uses a labelling technique which permits that the operator runs fast since it does not have to track edges separately. Typically about 5 seconds of cpu is required on a VAX 11/780 on a 256x256 image.

The operator is found to perform relatively successfully on the majority of the Alvey car images. One problem with of the operator is caused by the fact that it operates globally over a connected edges in the image. If a smooth man-made edge merges into a wiggly natural edge the discrimination is dependent upon the relative proportions of man-made to natural components. Likewise if an edge line contains forks the algorithm makes no distinction between branches of the line. This latter problem may be overcome by removing the fork points of each edge line before running the algorithm.

3.2 Natural Scene Labelling Techniques

This section briefly describes the work detailed in (Fretwell, Bayliss, Radford, 1988).

A very major problem with many symbolic approaches to object recognition is the computational load they require increases (often exponentially) with the number of features which are present in the image. To minimise these problems it is desirable to have some cue which may be extracted from the image which will serve as an interest operator for the high level algorithm.

The fractal discriminant described above provides one method for cueing to man made objects. Even after it's application however a large number of objects may be seen in the scene. For the problem domain of interest to the consortia (ie recognition of cars in urban scenes) it was found advantageous to consider the dominant edge direction in areas of the image. After the contorted lines have been masked from the image by use of the fractal discriminant described above, the majority of lines remaining in the image are found to originate from roads, houses and road vehicles. It was recognised that a useful cue to the identity of the object could be gained by analysis of the statistical distribution of edge directions within regions of the image. Thus the dominant emphasis of buildings is vertical, while cars have more horizontal lines. A simple routine was developed to analyse the regions of the image in which the distributions occurred and to mask regions not likely to contain a car. The method provided effective and, for most of the test images, proved capable of correctly rejecting approximately 60% of the image area as not containing cars. This greatly assists subsequent operations such as the Dynamic Programming shape detection technique of RSRE (Sleigh,1988) or other higher level operations.

3.3 Application Independent Theoretical Studies

3.3.1 Automated Inference Work

This section describes the work reported in (Fretwell, 1985a).

A major problem with the use of logic programming languages is the comparatively inefficient search mechanism. Since computational load incurred through the use of PROLOG for reasoning about image features is often very large, it is desirable to be able to guide the search path, either through incorporation of some heuristic derived from prior knowledge, or by an adaptive mechanism which monitors the paths which lead to successful search, and directs subsequent searches along the more successful routes first. In this work package a method has been developed to provide an adaptive search mechanism which operates within the normal PROLOG language.

To modify the search mechanism we adopt an approach in which the knowledge is contained in an evaluation function defined on the nodes of the search space. This function is a measure of how often each node has been involved in a minimal length solution tree. It is used to change the way the nodes are searched in the PROLOG program. The aim is to focus the deduction on parts of the search space which have, in the past, led to short solution paths. Since the quality of the search may be defined at either a global or a local scale several possible options are available.

The first strategy chosen for investigation uses the evaluation function to produce a search strategy that was a combination of breadth first and localised depth first search. The local depth first search is activated by the value of the function associated with the node exceeding some predetermined value. The search continues depth first provided the evaluation function values of subsequent nodes are above this threshold. This strategy has the benefit of focusing the deduction on parts of the search space which have in the past led to short solution paths. It will therefore save the deduction from making an extensive search, through possibly the whole space. This is a major consideration when the space is very large, for example a database consisting of logical descriptions of objects for an image understanding system.

The second method uses the evaluation function values in a more straightforward manner. The order of the nodes in the original program is altered by rearranging according to the evaluation function values associated with the nodes. The node with the highest evaluation function value appears at the top of the program. Thus the nodes with the highest utility are selected first by the standard prolog depth first search. This will encourage restrict the search towards more profitable areas of the search tree.

Experimental results have indicated that the two control methods have a part to play in developing a search control method which learns to adapt and improve search time in the prolog deduction mechanism. The main disadvantage the system faced was that it became saturated after prolonged learning. This was attributed to the simple learning method. However, if a more sophisticated description was employed some of these problems may be overcome.

Future work on control strategies for PROLOG could include topics such as consideration of deduction complexity and efficiency. For example the approach adopted at RSRE does not consider the time taken to perform the operations on the solution path. The simple philosophy adopted here of assigning worth or usefulness to a minimum solution path will not always not produce a minimum execution time for the goal. A much more sophisticated approach to the problem has been explored at Reading where an assumption based truth maintenance system has been developed.

3.3.2 Essence Description

The work outlined in this section is described more fully in (Fretwell, Goillau, 1988).

The long term objective of the Essence description is to develop tools and techniques for recognition in unconstrained scenes. The main thesis is that unconstrained image interpretation requires novel information structures and processing techniques. In particular, it is proposed to demonstrate that in principle a suitable knowledge formalism can be used to represent spatial and non-spatial information for scene interpretation, and secondly to show how this knowledge can be organised and used for recognition.

The general framework is one in which spatial information is supplemented with information about other object attributes to form a model. An example of the type of object attribute is the objects function or the context in which it usually appears. The relationship between an objects function and the resulting appearance of the objects has great potential in a representation system but the implementation of such ideas presents major problems.

Incorporation of functional attributes as well as purely spatial information in the object database can prove advantageous in a number of scenario's.

Most obviously, within the aims of the consortia, it is desirable to be able to describe objects in terms of their class. If the object description is limited to purely spatial attributes, then severe limitations may be placed on the number of manifestations which may be described by a single model. Although it is always possible to increase the range of objects which may be described as members of a given class through the inclusion of additional models, such an approach will greatly enlarge the database and consequently excassibate the matching process. It is possible that inclusion of functional or other non-spatial information may circumvent this problem. A second important advantage of the inclusion of non-spatial information is that it can greatly enhance the potential for the use of contextual or non-visual information. Often a vision sub-system will form only a part of an overall system which integrates a large number of sensors. input of collateral information from such sensors may prove of value in formulating and testing of object hypotheses in the vision sub-system.

During the project a methodology has been developed which takes a stylised description of an object in a stylised natural language form and builds a PROLOG description of the object. The description comprises an account of the salient features of the object, the terminal symbols forming hypotheses which may be tested by low level operators against the image data, or from inputs from other sensor sub-systems. Although, in the general case the translation takes place by hand, some progress has been made towards automating the process.

In hand-crafting human object descriptions to PROLOG, the system designer applies his own semantic processing which is at present difficult to quantify to extract key concepts and their relationships. With automatic translation of English descriptions there is a problem when mapping the imprecision and vagueness inherent in natural language descriptions (e.g. "sometimes", "mostly") onto the bimodal logic, negation by failure operation of PROLOG.

At present the method proposed matches the model against object features. If the background were to become cluttered, or more parts were added to the scene, then the matching mechanism would be affected adversely. The number of combinations of features necessary to identify the object would become expontentially large as the background and other parts in the image contributed more and more features that did not belong to the object. In general, open world situations provide uncertain and inconsistent data. Thus PROLOG with it's closed world reasoning is not well fitted to the task of relating the model to the image.

A deeper understanding of representation in image understanding problems was gained by using function. The essence work lead to the conclusion that the representation of image uncertainty (either because of the objects variability or poor image operators acting on the image) was crucial in the essence representation. Without it high level representations could not hope to communicate with the data. Thus the research was clearly pointed in the direction of obtaining more reliable image derived data and was of representing the uncertainty inherent in image information.

3.3.3 Image Verification using FRIL

A significant result from the study of the PROLOG approach to object description has been an appreciation of the constraints forced on the system by the adoption of a closed world bimodal logic. In an attempt to alleviate some of these problems a study was made of the value of a fuzzy logic programming language.

Through links with Bristol University fostered by the consortia it was known that a fuzzy symbolic language FRIL had been developed. At RSRE work was initiated to study the potential of FRIL for object description while Bristol University concentrated on development of Fuzzy conceptual graph methods. In addition RSRE funded additional effort at Bristol University to develop machine learning techniques based on conceptual graph descriptions.

The language, FRIL, permits the expression of incomplete or uncertain information and provides a calculus for combining evidence according to several models based on probability theory. The ability to express incomplete or uncertain knowledge provides a framework better suited to open world situations. This language was selected as being one of the best probabilistic inference languages available and through interaction with Bristol University through the MMI007 consortium.

The image is represented as a set of image features and two dimensional spatial relationships between them. The terminal nodes represent the output delivered by low level operators from the image. In contrast to the PROLOG system, the feature detectors return the confidence interval representing the probability that the feature has been detected (and possibly of it's absence). The object database is chosen to be typical of the relationships which might be expected to be found in real images that contain the object. Verification of the target object is obtained by determining those features and their relationships found in an image and matching them against the model.

The inference process in FRIL is used to effect the matching against image derived features. FRIL uses Zadeh's Fuzzy set theory and the Dempster-Shafer theory of evidence specialised to apply to the particular form of knowledge representation and inference mechanism of the support logic programming system.

The features are in general very variable in shape and will not consistently appear when the object is present in the image. This lack of certainty can arise from obscuration of object under consideration, poor algorithms for detecting image features such as edges, objects being near the limit of imaging resolution and so forth. In general it is possible to design a feature detector which will return a measure of confidence for both the presence or absence of each feature.

In practice feature detectors will not behave independently. For example variations in scene illumination may cause all detectors top become unreliable. Alternatively, in one type of scene it may be common for features close to ground level to become obscured while those higher up remain reliable. View angle is ano'her common problem. Some groupings of lines will be less distorted with view angle than others.

In the case of independent feature detectors the FRIL description of an object is similar to the PROLOG description. More computation is however required to evaluate the evidence, not just because of the algorithms for dealing with the uncertainty, but because it is necessary to evaluate all possible paths. Thus the heuristics to speed the search which were developed for PROLOG are no longer useful if the full solution is required. A more serious problem arises in the case of dependent feature detectors. Here a separate rule is required to describe the dependency between the feature detectors. Thus not only is the computational load greatly increased, but it is also necessary to have a very much more elaborate object model.

During the project it became clear that it was necessary to be able to evaluate the performance of feature detectors on a large number of images, particularly to determine the conditional dependence between detectors. Unfortunately insufficient time remained in the project to be able to develop and extract sufficient information for a detailed analysis of the situation.

4 Conclusions

The collaboration has provided a valuable research experience. Prime advantages of the collaboration arise from the increased access to the work of other sites. This arises both from the formal exchanges of reports and technical meetings, and also from the informal interactions which have arisen during the collaboration. The collaboration has permitted more work to be covered than could have been achieved had all sites been working in isolation.

There have been prices to pay. One of the greatest is in the loss of flexibility of the internal programme. Once the collaboration has been started, changes to the programme which must inevitably occur must be discussed between the partners and this can limit the ability to respond to local pressures (sometimes this may be played to advantage as well !). Other disadvantages are the considerable time taken by the overheads of the collaboration, both in setting up the consortia and also in porting software to unfamiliar machines.

One of the main managerial lessons learnt from the collaboration is that many of the disadvantages are of a fixed nature while most of the advantages scale with the level of involvement. Thus collaboration is only really of value if the level of involvement is more than one man. As a result collaborations are of most value to large centres. Clearly this is a major disadvantage in areas such as image processing where the number of experienced personnel are quite small and it is highly counterproductive to have them engaged in the many non-technical negotiations which collaboration incurs.

Technically the majority of the work carried out at RSRE under this collaboration has been involved in applications of symbolic approaches to both high and low level vision.

Two issues have been raised. First the computational complexity of the symbolic methodologies is a serious issue when more than a comparatively small number of features are being studied. This seriously limits the performance of the system in open world situations. When using a logic programming language such as PROLOG it was found that the native search mechanism, a depth first search, was generally undesirable for the application in hand, and that effort was expended in restructuring the mechanism thus negating one of the languages main features.

A more serious issue which has been exposed is the relative fragility of the symbolic approach. As the representation progresses to higher levels the effect of errors in earlier stages do not die out but have a tendency to propagate. Thus a small mistake in the early stages of processing may have a disastrous effect at the final stages. The adoption of a fuzzy logic does not remedy this problem, only the manifestation differs. It is found that uncertainty quickly dominates.

The lessons learned under this project about symbolic logic have had a large effect on the future direction of the section. At the outset of the project the main perceived problem was of representation of an expert's knowledge. With the benefit of the research carried out under this contract we now judge the dominant issue as one of formulating the problem as one amenable to efficient search and in a manner which is capable of being firmly based on probability theory. To this end work has been initiated to evaluate the applicability of a number of sub-symbolic tools, both as an intra-mural programme and as a collaborative IED proposal.

Principle achievements have been

- Development of a knowledge rich edge finder
- Development of a fast accurate edge finder suitable for SIMD and parallel machines.
- The use of fractal dimension for line classification has been demonstrated.
- The use of statistical measures for cueing has been researched
- PROLOG has been adapted to help improve performance when reasoning about image features.
- An appraisal of fuzzy logic for reasoning about uncertain features has been undertaken.

References

- BLAKE, A., (1983), *The Least-disturbance Principle and Weak Constraints*, Pattern Recognition Letters 1, 393-399.
- BOOTH, D., (1988), *Edge finder evaluation*, RSRE RIPR report 1000/17/88
- CANNEY, J.F., (1983), *Finding Edges and Lines in Images*, MIT Tech. Rep. AI-Tr-720.
- MARR, D., HILDRETH, E.,(1980), *Theory of Edge Detection*, Proc. Royal Soc. B 207, 187 -217.
- SLEIGH, A.C., (1988), *Segmentation and Concatenation of Edgel Lists by Dynamic Programming and Stochastic Models*, RSRE SP4 Research Note, Number 56, June, 1988.
- WITKIN, A., (1983), *Scale Space Filtering* , Proc. 7th IJCAI, 1019 - 1021.

5 RSRE Bibliography

FRETWELL, P., (1985a), RSRE, *A report on the Alvey Vision and Image Interpretation meeting 15-18 September, 1985*, RSRE SP4 Research Note 015, October 1985.

FRETWELL, P., (1985b),RSRE, *Search Control Strategies for PROLOG*, RSRE Research Note No. 11.

FRETWELL, P., GOILLAU, P.J., AND HEARN, D.B., (1987), RSRE, *Using the Essence of an object in computer recognition.* RSRE SP4 Research Note 28, March 1987.

FRETWELL, P., GOILLAU, P.J., HEARN, D.B., AND SLEIGH, A.C., (1987), *Representation and reasoning using linguistic descriptions of objects.* Paper presented at BPRA Colloquium 'Representing knowledge in vision', Imperial College, London, April 7, 1987.

FRETWELL, P., GOILLAU, P.J., HEARN, D.B., AND SLEIGH, A.C., (1987), *Linguistically defined generic models in computer vision.* Poster presented at Alvey Conference, UMIST, Manchester, July 14-16, 1987.

SLEIGH, A.C., HEARN, D.B. AND FRETWELL, P., (1986), *Can the function of an object be used in computer recognition?*, Paper presented at IEE Colloquium on "Knowledge based image processing", Savoy Place, London, January 27, 1986.

SLEIGH, A.C., (1986), *The extraction of boundaries using local measures driven by rules.* Pattern Recognition Letters 4, 247-258, 1986.

GOILLAU, P.J, (1987), RSRE, *Language interface and dialogues* , 1987, SP4 Research Note, No, 35. 3 July 1987.

FRETWELL, P., BAYLISS, D., RADFORD, C.J., (1988),RSRE, *Generic cueing in image understanding*, RSRE Memo. No. 4109.

FRETWELL, P., GOILLAU, P.J., (1988), *Linguistic definition of generic models in computer vision*, pp 306-315, in Lecture Notes in Computer Science, J. Kittler (ed), Pattern Recognition 4th Int. Conference Cambridge, UK, March 1988, Proceedings. Published by Springer Verlag

FRETWELL, P., BAYLISS, D., RADFORD, C.J., (1988), *Generic cueing in image understanding*, pp 325 - 333 in Lecture Notes in Computer Science, J. Kittler (ed), Pattern Recognition 4th Int. Conference Cambridge, UK, March 1988, Proceedings. Published by Springer Verlag

SLEIGH, A., (1986), *Extraction Boundaries using Local Measures Driven by Rules*, in Pattern Recognition letters, Vol. 4, pp 247 - 258, 1986. (AOI/TR/RSRE/850601)

RADFORD, C.J., (1989), *Vehicle Detection in Open World Scenes Using a Hough Transform Technique*, IEE International Conference on Image Processing and Applications, July, 1989, Warwick, Great Britain.(To be published)

DOCUMENT CONTROL SHEET

Overall security classification of sheet ...UNCLASSIFIED...

(As far as possible this sheet should contain only unclassified information. If it is necessary to enter classified information, the box concerned must be marked to indicate the classification eg (R) (C) or (S))

1. DRIC Reference (if known)	2. Originator's Reference MEMO 4277	3. Agency Reference	4. Report Security Classification U/C
5. Originator's Code (if known) 7784000	6. Originator (Corporate Author) Name and Location ROYAL SIGNALS & RADAR ESTABLISHMENT ST ANDREWS ROAD, GREAT MALVERN, WORCESTERSHIRE WR14 3PS		
5a. Sponsoring Agency's Code (if known)	6a. Sponsoring Agency (Contract Authority) Name and Location		

7. Title
Alvey MMI007: Object identification from 2D images...

7a. Title in Foreign Language (in the case of translations)

7b. Presented at (for conference papers) Title, place and date of conference

8. Author 1 Surname, initials	9(a) Author 2	9(b) Authors 3,4...	10. Date	pp. ref.
FRETWELL P	SERIES, R		1989.05	17
11. Contract Number	12. Period	13. Project	14. Other Reference	

15. Distribution statement
UNLIMITED

Descriptors (or keywords)

continue on separate piece of paper

Abstract

A consortia was formed in 1984 to tackle the subject of automatically finding and identifying specified objects in a complicated environment from two dimensional images. The consortia members consist of four industrial partners - British Aerospace, Marconi Command and Control Systems, Standard Telecommunications Laboratories and The Royal Signals and Radar Establishment - and three universities - Bristol, Reading and Surrey. The project began in 1985 and now in 1989 is almost completed. This report describes the RSRE contribution to the MMI007 project.

S80/48